

AMENDMENTS TO THE SPECIFICATION

Please amend paragraph [0031] as indicated below:

[0031] Figure 3 is a virtualizer architecture with JBOD and RAID 300 that includes the elements of virtualizer architecture with RAID 200 as well as a JBOD 310 coupled to virtualizer module 1 220. Virtualizer architecture with JBOD and RAID 300 allows for coalesced writes to JBOD 310. A coalesced write is simply the process of collecting multiple write requests to a group of sequential or nearly sequential logical block addresses (LBAs) so that the data may be written with a single write command to sequential LBAs. This process minimizes tracking and seeking motions performed by the head which, in turn, minimizes the time required to perform the writes as well as minimizing the physical head motion. Minimizing head motion increases the longevity of JBOD 310 and thus increases the mean time between failures (MTBF). The following is an example of a coalesced write. The example is used for illustrative purposes only and in no way limits the actual implementation of virtualizer architecture with JBOD and RAID 300. In this example, host 1 210 issues a write command to an LBA residing on JBOD 310. Virtualizer module 1 220 receives the command and data, and stores the write data in its own cache and in the cache of virtualizer module 2 230. Virtualizer module 220 then sends a write acknowledge back to host 1 210. Host 1 210 issues a read command from an address on storage element 295. Next, host 1 210 issues another write command to the next sequential LBA residing on JBOD 310. Virtualizer module 1 220 also stores this data in its own cache and in the cache of virtualizer module 2 230 and sends an acknowledge back to host 1 210. Host 1 210 then performs a write to storage element 295. Finally, host 1 210 sends a third write command to JBOD 310 via virtualizer module 1 220. This command and data are also stored in cache, and virtualizer module 220 acknowledges the command to host 1 210. The cache of virtualizer module 1 220 now holds the data for three write commands that are to be

written to three consecutive LBAs on JBOD 310. Virtualizer module 1 220 creates a single write command from the three original write commands and sends the command and data to JBOD 310. JBOD 310 performs the three writes as a single write command and sends the complete acknowledgement to virtualizer module 1 220. The result is not only less wear and tear on the head of JBOD 310 but also in a reduction in latency. Using virtualizer architecture with JBOD and RAID 300, JBOD 310 finds the beginning LBA using a seek operation and performs the write for all three write requests. In a traditional system, JBOD 310 would need to locate three different LBAs, and then write three separate sets of data using multiple disk accesses at separate times.